Yarrp'ing the Internet: Randomized High-Speed Active Topology Discovery

Robert Beverly
Naval Postgraduate School
rbeverly@nps.edu

arXiv:1605.03999v1 [cs.NI] 12 May 2016

ABSTRACT

Obtaining a "snapshot" of the Internet topology remains an elusive task. Existing active topology discovery techniques and production systems require significant probing time – time during which the underlying network may change or experience short-lived dynamics. In this work, we consider how active probing can gather the Internet topology in *minutes* rather than days. Conventional approaches to active topology mapping face two primary speed and scale impediments: i) per-trace state maintenance; and ii) a low-degree of parallelism. Based on this observation, we develop Yarrp (Yelling at Random Routers Progressively), a new traceroute technique designed for high-rate, Internet-scale probing. Yarrp is stateless, reconstituting all necessary information from ICMP replies as they arrive asynchronously. To avoid overloading routers or links with probe traffic, Yarrp randomly permutes an input $IP \times TTL$ space. We run Yarrp at $100Kpps$, a rate at which the paths to all /24's on the IPv4 Internet can be mapped in approximately one hour from a single vantage point. We compare Yarrp's topological recall and discovery rate against existing systems, and present some of the first results of topological dynamics exposed via the high sampling rates Yarrp enables.

1. INTRODUCTION

Network and security researchers rely on topological maps of the logical Internet to address problems ranging from critical infrastructure protection to policy. Production active measurement systems that continually gather and curate Internet topology, e.g. [21, 15], are thus important to many longitudinal analyses and to shedding light on network events of interest.

Obtaining IP, router, and provider-level network topologies has been a continual research focus for more than two decades [14, 17, 25, 21]. While significant progress has been made, topology mapping at Internet scale remains a challenge. Both the accuracy of the inferred network topologies [27], and the speed at which they can be recovered [10, 6], present obstacles to current mapping efforts. In this work, we focus on the *speed* and *scale* of Internet-wide active topology mapping.

Given its scale, experience and popular belief dictates that obtaining even partial Internet topologies via active network probing is a time-intensive process. For instance, CAIDA's Archipelago (Ark) [15] system uses dozens of vantage points and at least a day to traceroute to a single address in each routed /24 IPv4 prefix. A recent topology cycle gathered by Ark from April, 2016 [8] sent approximately 11M traceroutes from 37 monitors over the course of 31 hours in order to discover ∼1M distinct router interfaces, and ∼2M links.

We re-examine some of the assumed fundamental limits of active topology mapping to consider whether such probing could be performed in minutes rather than hours. Taking inspiration from recent stateless and randomized high-speed network scanners such as ZMap [11] and masscan [12], we create Yarrp (Yelling at Random Routers Progressively).

To facilitate high-probing rates, Yarrp is stateless, reconstituting all necessary information from replies as they arrive asynchronously. To avoid overloading routers or links, Yarrp randomly permutes its input $IP \times TTL$ space when probing. Yarrp is thus well-suited for Internet-scale studies. Our contributions include:

1. Development of Yarrp, a publicly available tool [5] that permits rapid active network topology discovery. We run Yarrp at $100Kpps$ to discover more than 400,000 router interfaces in under 30 minutes.
2. A comparison of Yarrp and CAIDA's existing production topology collection platform, showing recall and speed differences.
3. As an application of rapid topology discovery, we conduct successive topology snapshots separated by a small time delta and characterize the distribution and causes of observed path differences.

2. BACKGROUND

Traditional traceroute [16] obtains the sequence of router interface IP addresses along the forward path to a destination by sending probe packets with varying time to live (TTL) values and examining the ICMP responses. By maintaining the transmission timestamp of each probe, traceroute can report the round trip time (RTT) from the source to each responsive hop. Mod-

ern traceroute implementations send batches of concurrent probes to lower tracing time, e.g. Linux defaults to 16 simultaneous probes. In order to match probes to the ICMP TTL exceeded responses they generate, the probe must include unique identifiers that are returned as part of the ICMP quotation [22]. Because the quote is only required to copy the first 28 bytes of the packet that induced the expiry message [3], traceroute typically relies on the first 8 bytes of the transport-layer header to match responses to probes.

While various improvements have been proposed and implemented, the core behavior of traceroute – and large scale active topology scanning – remains largely unchanged. To prevent false inferences due to load-balanced paths, Augustin et al. created Paris traceroute [2]. To reduce unnecessary probing, Donnet et al. developed Doubletree [10], a modified traceroute that begins probing from a likely path midpoint outward until it reaches previously discovered hops. Similarly, [6, 4] proposed several topology primitives empirically shown to reduce the volume of probing while maintaining or increasing topological discovery. Luckie et al. developed Scamper [20], a production-quality packet prober. Scamper implements both Doubletree and Paris traceroute, has an open API, can maintain a configurable probing rate, and can be controlled remotely. CAIDA's production Ark infrastructure [15] uses Scamper to perform continual Internet-wide probing [8].

Traceroute was originally designed as a tool for network administrators to diagnose a small number of paths, not as a means to gather snapshots of the entire Internet topology [16, 18]. Fundamentally, traceroute and its variants all have two properties that limit their scalability and speed. They:

- Maintain state for each outstanding probe sent, including some identifier and origination time.
- Are sequential, probing all hops along the path to a destination. While some tools (e.g. scamper) can traceroute to multiple targets, this parallelism is limited to a finite window of destinations.

In contrast, Yarrp is designed to be stateless and random – probing different portions of many different paths simultaneously. This allows Yarrp to send probes at a high per-packet rate, while spreading the load among many destination networks to avoid concentrating load on particular paths, links or routers, thereby avoiding anomaly alarms or ICMP rate limiting.

3. YARRP DESIGN

The high-level idea of Yarrp is: i) randomization of the probing order of the domain of network range(s) and TTLs; and ii) stateless operation, whereby all necessary state is encoded into the probes such that it can be recovered from the ICMP replies. Yarrp is written in C++, is portable to a variety of UNIX-like platforms, and is publicly available [5].

3.1 Pseudo-random Probing Order

Existing traceroute techniques probe all hops along a path to a destination in sequence. Instead, we employ a keyed block cipher to provide a bijection over the input domain of target IPs and TTLs ($D = IPs \times TTLs$). This means that Yarrp will e.g. send a probe to IP address A with $TTL = 10$, then toward B with $TTL = 3$, then C with $TTL = 7$, and so on until the entire space of $TTLs$ for each target IP is covered.

The symmetric RC5 block cipher with a 32-bit block size is fast and a natural fit for our application[1]. With key k, Yarrp encrypts the sequence $i = 0, 1, \ldots, |D| - 1$ where bits of each ciphertext $C_i = RC5_k(i)$ determine the target IP address and TTL to probe. In this way, Yarrp randomizes the order of probed $\langle targetIP, TTL \rangle$. Yarrp can permute arbitrarily large or small IPv4 address and TTL domains, or can permute the order of specific targets read from a file. Depending on the size of the domain, we switch between either a prefix-cipher or cycle-walking cipher, as described in [7].

To facilitate comparison with CAIDA's IPv4 topology dataset [8], Yarrp has a mode that probes a random address in each IPv4 24-bit subnet – this mimics the targets selected in a full cycle of CAIDA's probing. Here, Yarrp encrypts each $i = 0, \ldots, 2^{24} - 1$ with key k. For $C_i = RC5_k(i)$, Yarrp probes the IPv4 address $C_i[0:23] * 2^8 + (C_i[0] + C_i[1] + C_i[2])\%256$ with TTL $C_i[24{:}31]$. In this fashion, we permute through the space of 2^{24} possible /24s, and construct the least-significant octet as a function of the subnet such that the same random address in each /24 is used as the destination for each TTL.

An advantage of Yarrp's randomization method is that the probing work can easily be distributed among multiple vantage points with negligible coordination or communication overhead. We discuss distributed Yarrp as a future enhancement in §5.

3.2 Stateless Operation

Existing traceroute techniques require state to match ICMP replies to probes. In contrast, Yarrp does not require state. We overload various fields in the probe packets with specific values such that we can reconstruct the corresponding probe's destination, transmission time, and originating TTL from within the quote of the ICMP TTL exceeded messages.

Figure 1 depicts the TCP/IP header fields we utilize. We encode the TTL with which the packet was sent in the IPID and the elapsed time in the TCP sequence number. We use elapsed time rather than e.g. Unix time

[1]Other block ciphers could be used; the cryptographic strength of the cipher is not critical to our application.

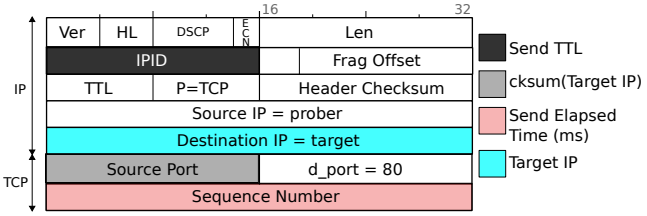

Figure 1: Yarrp encodes information in the IP and TCP fields of outgoing probe packets in order to permit stateless operation.

in order to maintain millisecond resolution with only a 32-bit field. Yarrp can also encode microsecond resolution, so long as the expected duration of a probing run is less than $2^{32}/1 \times 10^6 \simeq 4300$ seconds. The destination TCP port is fixed to port 80 to facilitate firewall traversal, while we populate the source TCP port with the checksum of the target IP address. In this fashion, we can detect instances where the destination IP address is modified enroute, a phenomenon Malone and Luckie observe in 2% of their results [22].

In order to properly accommodate load-balanced paths, which are common in the Internet, we ensure that, for a given destination, certain fields remain fixed for all TTL probe values. For instance, although the TCP source port changes, it is a function of the destination IP address and therefore will contain the same value for all probes sent toward the destination. This design allows us to maintain the benefits of Paris traceroute [2].

When ICMP TTL exceeded messages arrive, we examine the included quotation to recover the destination probed, the originating TTL (hop), responding interface at that hop, and compute the RTT by taking the difference of the packet arrival time and the probe origination time as encoded in the quoted TCP sequence number. These values can be computed from the minimum 28 bytes of required quotation [3].

Finally, Yarrp can source either TCP SYN or ACK probes. While SYN probes can permit middlebox traversal, we use the ACK-only mode to avoid alarms triggered by large volumes of SYN traffic.

3.3 Challenges

The benefits of Yarrp's design come with several concomitant challenges, namely: i) reconstructing the unordered responses into paths, ii) knowing when to stop probing, and iii) avoiding unnecessary probing.

In following with Yarrp's stateless nature, ICMP responses are decoded as they arrive and written sequentially to a structured output file. Each entry in the output file corresponds to an ICMP response. An entry includes the target IP address, originating TTL, responding router interface IP address, RTT, and meta-data such as timestamps, IPID, response TTL, packet sizes, and DSCP markings. Because of the inherently random probing, the entries for each hop along a path to a given

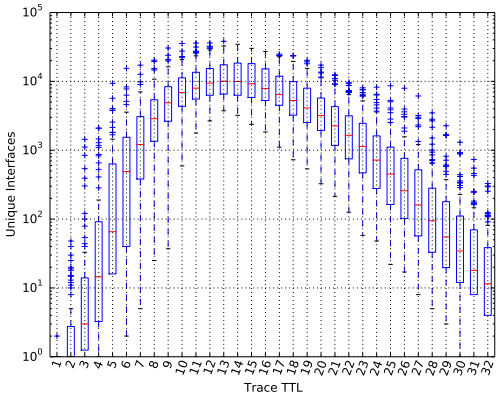

Figure 2: Distribution of interfaces discovered across Ark vantage points in a single cycle as a function of TTL (path depth). The red line indicates the median value among vantage points.

destination will be unordered and intermixed with other responses in the Yarrp output file. We must therefore reconstruct complete paths by parsing the entire output file and maintaining state for each destination. While this is a memory and time-intensive task, the key point is that it can be performed *off-line*. In this fashion, we decouple probing from path reconstruction to permit the probing to be as fast as possible. Included in the Yarrp distribution [5] is a `yrp2warts` Python script that performs this off-line conversion into the standard warts [20] binary trace format.

A practical consequence of Yarrp's randomization and lack of state is that its probing behavior does not depend on the received responses. Thus, Yarrp cannot stop probing once it reaches the destination or when the path contains a sequence of unresponsive hops (the so-called "gap limit"). To better understand the optimal range of TTLs to probe (from the possible space 1-255), we examine the results from a complete cycle of Ark probing from January, 2016 [8]. We seek to determine, across each of the Ark vantage points, the number of unique router interfaces discovered at each TTL.

Figure 2 shows the inter-quartile range of the number of distinct interfaces found as a function of TTL for each of the vantage points; the red line in the boxplot displays the median number of interfaces per TTL among each vantage point. Because of the Internet's tree-like structure, the first few hops reveal only a small number of interfaces regardless of the destination probed. The bulk of the interfaces are found between TTLs 10 to 16, with an inflection point around a TTL of 14. The amount of discoverable topology beyond a TTL of 32 is negligible (note the log scale y-axis). As a result, Yarrp defaults to probing TTLs 1 to 32 to minimize unnecessary probing while exploring the majority of the space.

For many destinations Yarrp will perform more probing than traditional traceroute methods. This is both

an advantage and a disadvantage: we show in §4.1 that discoverable topology exists beyond multiple unresponsive hops (where existing methods terminate early).

3.4 Optimizations

In environments sensitive to probing volume, several optimizations can substantially decrease unnecessary probing at the expense of maintaining some state. This subsection discusses optimizations to the base Yarrp design to enable different tradeoffs.

First, Yarrp can read a BGP routing table of network prefixes and build a longest-match Patricia trie [24]. When iterating through the entire permuted IPv4 space, Yarrp can skip destinations that are not routed. Based on current global BGP routing tables [1], this optimization avoids probing approximately 1.5B IP addresses (35% of the 32-bit space) that are unlikely to return useful results. Note that the memory required to maintain the BGP table is constant during a probing run (amounting to approximately 300MB during runtime). In our experiments, these lookups in the Patricia trie did not prevent Yarrp from running at over $100kpps$.

Second, the tree-like structure of the network implies that the set of interfaces near to the vantage point is small relative to the universe of router interfaces [10]. In Figure 2 for instance, all of the traces have a single first hop in common and orders of magnitude fewer interfaces at hops 1-3 as compared to hops 13-15. To avoid rediscovering the same nearby router interfaces repeatedly, Yarrp can maintain state over the set of responding local "neighborhood" interface IP addresses at hops 1 through a run-time configurable ttl_{nbrhd}. For each TTL in the neighborhood, Yarrp maintains two timestamps: the last time a probe was sent with that TTL, and the last time a new interface at that depth replied. If no new interfaces have been discovered within the past $\eta = 30$ seconds of probing, Yarrp skips future probes at that TTL. The `yrp2warts` script can then stitch together these missing hops. While the amount of state in neighborhood mode can grown unbounded, in practice it is small for small ttl_{nbrhd}, while avoiding substantial over-probing.

4. RESULTS

This section examines results from running Yarrp on the Internet. We compare the topological recall against an existing production system and then analyze the discovery yield (i.e. the amount of new topology discovered over time). Finally, as an application of Yarrp's probing speed, we gather two successive snapshots of Internet topology separated in time by a small delta to reveal instances of short-lived network dynamics.

4.1 Topological Recall

We empirically verify Yarrp's correctness by evaluating it against CAIDA's use of scamper in the Ark infrastructure [15, 20]. CAIDA makes their topology traces publicly available [8]. We find all 67,045 destinations probed from the San Diego vantage point on May 1, 2016. From this same network attachment point, we instruct Yarrp to probe the same destinations. In this fashion, because the vantage point networks and set of destinations are the same, we expect the paths to be largely congruent, allowing an unbiased comparison of topological recall between the two probing methods.

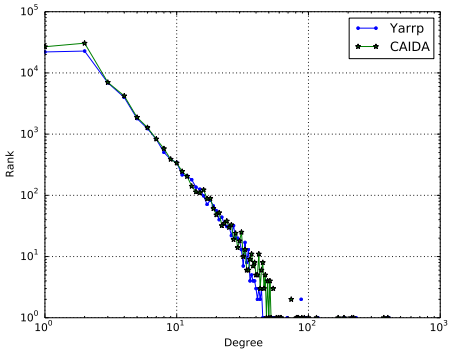

Figure 3: Degree distribution of topology gathered by Yarrp versus Ark when probing the same targets from the same source network.

From the Yarrp and Ark probing, we construct graphs of interface nodes connected by edges when the interfaces appear in consecutive hops of a path. We ignore anonymous interfaces [13] such that the graph may be disconnected. Figure 3 displays the resulting graph degree distributions on a log-log scale. We observe that the distributions match closely, empirically supporting Yarrp's ability to discover the responsive topology.

However, Yarrp discovers 16% fewer interfaces (80,134 versus 66,939) and 13% fewer edges (96,763 versus 84,113) than Ark in this experiment. We attribute much of this difference to Yarrp's use of TCP ACK probes as compared to CAIDA's use of ICMP probes. In a survey by Key et al., ICMP probes elicit almost 18% more responses as compared to TCP probes [19]. ICMP and UDP-based probing are future Yarrp enhancements (see §5).

Yarrp's stateless nature implies that it probes all TTLs from 1 to 32, whereas Ark's use of scamper ceases probing after encountering five unresponsive hops in a row. In the same CAIDA San Diego probing run, 39,613 traces stopped due to this gap-limit. For each of these gap-limit traces, we compute the difference of the highest responding TTL hop from Yarrp probing and the highest responding TTL from the Ark probing. Figure 4 shows the cumulative distribution of this difference among the gapped traces; a positive difference means that Yarrp discovered topology beyond the point where Ark stopped probing. For ≈74% of the traces, there is no difference. In 8% of the traces, Ark discovers one more hop than Yarrp. However, Yarrp discovers

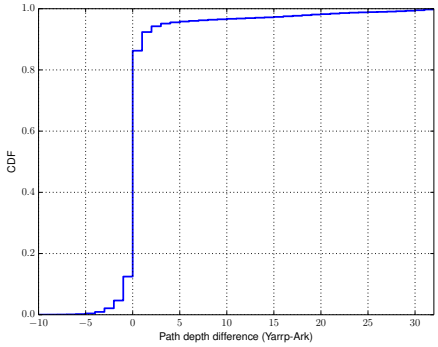

Figure 4: CDF of highest responding TTL difference (max(Yarrp)-max(Ark)) for Ark gap-limited traces (same dests. and vantage point).

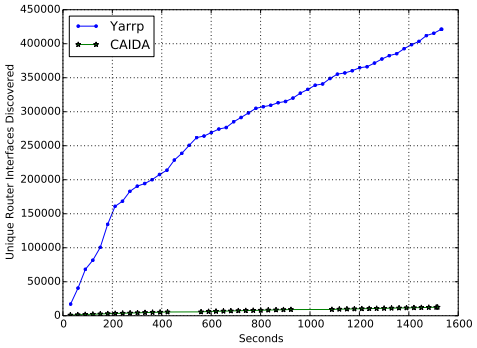

Figure 5: Unique router interface discovery rate: comparing Yarrp against CAIDA's routed /24 probing from a single vantage point.

one additional hop in $\sim$ 6% of the targets, and more than 5 additional hops in 4% of the cases.

4.2 Discovery Rate

A goal of Yarrp is rapid topological discovery. In this subsection, we look specifically at the ability to discover unique router interface addresses rapidly.

On May 10, 2016, we run Yarrp from a Northeast United States university vantage point at $\sim$ 100*kpps* and instruct it to perform the Ark-mode randomized probing of the globally routed IPv4 /24 prefixes. We limit Yarrp to this rate, and limit the duration of our experiment, per prior agreement with the local network administrator. The physical machine is a multi-core Intel L5640 processor running at 2.27GHz, with Yarrp running on an Ubuntu virtual machine allocated a single core. At this rate, the CPU utilization is $\approx$52%.

We enable the "neighborhood" optimization, as described in §3.4, as we are interested in finding as many distinct router interfaces as possible given the probing rate. Figure 5 displays the cumulative number of distinct router interfaces discovered as a function of time. As a basis of comparison, we also plot the number of unique interfaces found over time for a single vantage point (again, using data from the San Diego node of CAIDA's continual /24 probing [8] on May 1, 2016).

CAIDA's San Diego monitor discovers 12,568 unique interfaces in 1,500 seconds ($\sim$ 8 per second). By contrast, Yarrp discovers 421,162 unique IPv4 interfaces in the same period, or approximately 280 distinct router interfaces a second. The interfaces found by Yarrp in less than 30 minutes equates to 42% of *all* unique interfaces discovered from all Ark monitors over the course of probing for more than a day.

4.3 Short-Lived Dynamics

As an application of rapid topology discovery, we collect topology snapshots in rapid succession and analyze their properties and differences in this subsection.

We gather 67,045 target destinations from CAIDA's May 1, 2016 topology probing from their San Diego monitor. Again using the east coast university vantage point, we run Yarrp to probe TTLs 1-32 for these same 67,045 targets. We run Yarrp at $\sim$ 4000*pps* and invoke Yarrp three times in succession with a minute pause in-between. In this way, each snapshot takes approximately 8 minutes to gather, and each is separated by a minute. We term the snapshots S_1, S_2 and S_3 in chronological order.

The interface-level graph resulting from S_1 contains 39,968 interfaces and 46,721 edges, while S_2 has 40,038 interfaces and 46,773 edges. S_3 contains 39,994 interfaces and 46,749 edges.

To better understand the differences between snapshots, we perform a per-target path comparison. For each target in S_i, we compare the discovered path in S_i to the discovered path to that same target in S_j. We use the Levenshtein edit distance to measure the per-target path differences between snapshots. The edit distance is the minimum number of edits (insertions, substitutions, or deletions of router interfaces).

Note that inter-snapshot differences are not attributable to per-flow load balancing as Yarrp keeps the packet header fields which are used for load balancing constant for the same destination between snapshots (§3.2).

Additionally, to better understand the types of path changes, we count the frequency of each edit operation and missing hop substitutions. These missing hop operations are instances where the path contains a responsive router for a particular TTL for one snapshot, but no response at that TTL when probing the same destination in a subsequent snapshot. Such missing hops may be attributable to routers performing ICMP rate limiting, or may be due to packet loss.

A deeper analysis of the most frequent missing hops between S_1 and S_2 reveals that the large majority (92.2%) come from the first four hops within the local network of the vantage point. Specifically, 73% of the missing hops are due to the router at TTL 3, 18% are due to the router at TTL 1, and 1% are due to the router at TTL

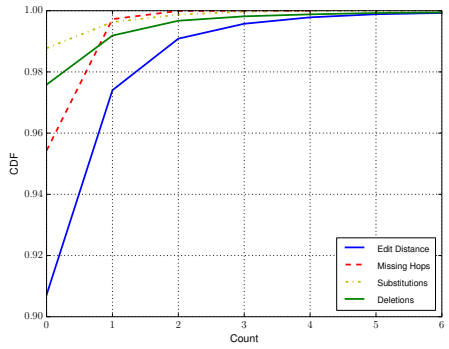

Figure 6: Cumulative fraction of measured paths versus pair-wise path edit distance between successive Yarrp topology snapshots.

4. In contrast, the router at TTL 2 always responds, suggest that some of the local routers implement ICMP rate limiting while one does not.

Figure 6 displays the results of the edit distance comparison between S_1 and S_2, ignoring differences attributable to the local network (TTL $\leq$ 4, as described above). The paths to approximately 91% of the destinations are identical between S_1 and S_2, while approximately 6% have a single hop difference. Less than 1% of the destinations show a difference of ≥ 2 hop edits.

Separated by the edit operation, we see that $\sim$ 4% of the 1 hop differences are due to missing hops, 1% are hop deletions, and fewer than 1% are substitutions.

To understand the potential of rapidly collected topology snapshots, we manually investigate and highlight a path exhibiting a significant change between snapshots. Figure 7 shows, for each of the three snapshots, the final four responsive hops toward the destination 188.32.230.138 (in AS 42610). In the intermediate snapshot, S_2, we see two hop substitutions where the next hop after AS 1273 (213.185.219.106, Cable and Wireless) changes to a different sequence of hops within AS 12389 (Rostelecom). By the time of the third snapshot, S_3, the path changes back to that seen in S_1.

While the exact cause of this short-lived routing change is unknown, the key point is that it would not have been discovered by the existing topology mapping systems. We leave a comprehensive analysis of the extent and duration of these short-lived dynamics exposed by Yarrp to future work.

5. CONCLUSIONS

Yarrp demonstrates a new technique for Internet-scale active topology probing that permits rapid collection of topology snapshots. As with our initial investigation of short-lived topology dynamics, our hope is that Yarrp facilitates analyses not previously possible.

Yarrp is stable and the code is publicly available [5]. That said, there are several enhancements that would

```
S1: 213.185.219.106 95.167.95.110 94.25.8.173 77.37.254.198
S2: 213.185.219.106 87.226.133.50 188.254.78.129 77.37.254.198
S3: 213.185.219.106 95.167.95.110 94.25.8.173 77.37.254.198
```

Figure 7: Example of short-lived path dynamics only observed via successive, rapidly collected topology snapshots.

be valuable additions. First, Yarrp currently only supports IPv4 probing. Given the vastly larger IPv6 address space, and relative topological sparsity [23], adding IPv6 support to Yarrp could enable more complete maps of the IPv6 topology to be gathered.

Second, Yarrp is currently capable of sending only TCP probes. It is well-known that using different transport protocols yields different responses, due to widespread security and policy filtering [19]. We plan to add ICMP and UDP probing to Yarrp, which requires utilizing different transport header fields to encode probe information. Doing so is non-trivial as we must maintain both the Paris-traceroute property of keeping certain fields constant to keep packets on a single load-balanced path, while also retaining Yarrp's stateless behavior.

Third, Yarrp's stateless and asynchronous nature implies that a malicious actor could attempt to send bogus responses, while middleboxes are known to mangle packet headers [9, 22]. In the future, we wish to use a keyed cryptographic integrity function over multiple probe values. Instead of a simple checksum on the target IP address, we will populate the source port with the value of this keyed integrity check. Yarrp can then ensure that it both sent the original probe, and that the probe was not modified in-flight such that the response is not useful.

Finally, an attractive feature of Yarrp's design is the ability to easily randomize and distribute the probing to multiple vantage points with negligible coordination and communication overhead. Similar to the rapid scanning worm envisioned by Staniford et al., the permuted domain can be distributed [26]. For n vantage points, and a given domain $D = IPs \times TTLs$, each vantage point encrypts $\frac{|D|}{n}$ of the range $i = 0, 1, \ldots, |D| - 1$. Thus, the v'th vantage point encrypts the range of values $\frac{|D|}{n}v$ to $\frac{|D|}{n}(v+1)$ to obtain its sequence of $\langle targetIP, TTL \rangle$ values to probe from the overall permutation. The potential speed improvement is linearly proportional then to the number of vantage points. Only the values $|D|$ and v need be sent to each vantage point to distribute the permuted space and achieve complete randomized coverage. Given our empirical (and conservative) $100kpps$ Yarrp rate in this work, we estimate that it is possible to implement a distributed Yarrp among $\sim$128 vantage points to traceroute to every routed IPv4 address ($\approx 2^{31}$ targets) in approximately one hour. Yarrp may thus facilitate rapid collection of *complete* Internet snapshots in the future.

6. REFERENCES

[1] University of Oregon RouteViews, 2016. `http://www.routeviews.org`.

[2] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira. Avoiding traceroute anomalies with Paris traceroute. pages 153–158, 2006.

[3] F. Baker. Requirements for IP Version 4 Routers. RFC 1812 (Proposed Standard), June 1995. Updated by RFCs 2644, 6633.

[4] G. Baltra, R. Beverly, and G. G. Xie. Ingress Point Spreading: A New Primitive for Adaptive Active Network Mapping. In *Proceedings of the 15th Conference on Passive and Active Network Measurement (PAM)*, volume 8362, pages 56–66, Mar. 2014.

[5] R. Beverly. Yarrp, 2016. `https://www.cmand.org/yarrp`.

[6] R. Beverly, A. Berger, and G. G. Xie. Primitives for active Internet topology mapping: toward high-frequency characterization. In *Proceedings of the ACM 10th annual conference on Internet measurement (IMC)*, pages 165–171, 2010.

[7] J. Black and P. Rogaway. Ciphers with arbitrary finite domains. In *Topics in Cryptology–CT-RSA*, pages 114–130. Springer, 2002.

[8] CAIDA. The CAIDA UCSD IPv4 Routed /24 Topology Dataset, 2016. `http://www.caida.org/data/active/ipv4_routed_24_topology_dataset.xml`.

[9] R. Craven, R. Beverly, and M. Allman. A Middlebox-cooperative TCP for a Non End-to-end Internet. In *Proceedings of ACM SIGCOMM*, pages 151–162, Aug. 2014.

[10] B. Donnet, P. Raoult, T. Friedman, and M. Crovella. Efficient algorithms for large-scale topology discovery. 33(1):327–338, 2005.

[11] Z. Durumeric, E. Wustrow, and J. A. Halderman. Zmap: Fast internet-wide scanning and its security applications. In *USENIX Security*, pages 605–620, 2013.

[12] R. Graham, P. McMillan, and D. Tentler. Mass Scanning the Internet. In *DEF CON 22*, 2014.

[13] M. H. Gunes and K. Sarac. Resolving anonymous routers in Internet topology measurement studies. *INFOCOM. The 27th Conference on Computer Communications*, pages 1076–1084, 2008.

[14] B. Huffaker, M. Fomenkov, and k. claffy. Internet topology data comparison. *Cooperative Association for Internet Data Analysis (CAIDA)*, 2012.

[15] Y. Hyun and k. claffy. Archipelago measurement infrastructure, 2015. `http://www.caida.org/projects/ark/`.

[16] V. Jacobson. traceroute, 1989. `ftp://ftp.ee.lbl.gov/traceroute.tar.gz`.

[17] k. claffy, Y. Hyun, K. Keys, and M. Fomenkov. Internet mapping: From art to science. In *IEEE Cybersecurity Applications and Technologies Conference*, Mar. 2009.

[18] k. claffy, Y. Hyun, K. Keys, and M. Fomenkov. Internet mapping: from art to science. In *Proceedings of IEEE Cybersecurity Applications and Technologies Conference for Homeland Security*, Mar. 2009.

[19] K. Keys, Y. Hyun, M. Luckie, and k. claffy. Internet-Scale IPv4 Alias Resolution with MIDAR. *IEEE/ACM Transactions on Networking*, 21(2):383–399, Apr 2013.

[20] M. Luckie. Scamper: a scalable and extensible packet prober for active measurement of the Internet. In *IMC*, pages 239–245, Nov. 2010.

[21] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: An information plane for distributed services. In *Proceedings of USENIX OSDI*, Nov. 2006.

[22] D. Malone and M. Luckie. Analysis of ICMP quotations. In *Proceedings of the 8th Passive and Active Measurement (PAM) Workshop*, Apr. 2007.

[23] J. Rohrer, B. LaFever, and R. Beverly. Empirical Study of Router IPv6 Interface Address Distributions. *IEEE Internet Computing, Special Issue on Measuring the Internet*, July 2016.

[24] K. Sklower. A tree-based packet routing table for berkeley unix. In *USENIX Winter*, volume 1991, pages 93–99, 1991.

[25] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP topologies with Rocketfuel. *ACM SIGCOMM Computer Communication Review*, 32(4):133–145, 2002.

[26] S. Staniford, V. Paxson, N. Weaver, et al. How to own the internet in your spare time. In *USENIX Security Symposium*, pages 149–167, 2002.

[27] W. Willinger, D. Alderson, and J. C. Doyle. Mathematics and the Internet: A source of enormous confusion and great potential. *Notices of the AMS*, 56(5), 2009.